

REMARKS/ARGUMENTS

Reconsideration of the present application, as amended, is respectfully requested.

The August 11, 2006 Office Action and the Examiner's comments have been carefully considered. In response, claims are amended and remarks are set forth below in a sincere effort to place the present application in form for allowance. The amendments are supported by the application as originally filed. Therefore, no new matter is added.

REQUIREMENT FOR INFORMATION

In the Office Action the Examiner states that there are differences between the present application and the International Application from which the present application claims priority. The Examiner also states that some of the claims do not appear to be supported by the provisional application, which is also claimed as a basis for priority. The Examiner requests that Applicant point out which portion of the disclosure in the provisional application and the International Application provide the written description and enablement support for each claim.

In response, Applicant states that support for each claim is as follows:

Claim 1

Provisional Application, written description support for Claim 1:

See page 1, lines 5 - 6:

The invention relates to the field of mirroring and in particular to asynchronous remote mirroring for Storage Area Networks (SANs).

See page 1, lines 14 - 24:

The description relates to the capability to replicate the contents of a storage device, pertaining to a local SAN (Storage Area Network), in part or in whole, to a remote-SAN coupled to the local SAN replication or mirroring understood as copying, are applicable to any object, volume, file, data structure or the like. The purpose of a replication is often to enable fast recovery from data loss caused by a disaster, either a fire, or a tornado or a hardware crash, etc. Remote mirroring means replication from an object located in one SAN to another SAN. Reference is made to the co-pending patent application PCT/IL00/00309, named "Storage Virtualization in a Storage Network," by the same applicant, incorporated herewith by reference in whole, and cited below as the '309 patent.

See page 4, line 1 to page 6, line 8:

Mirroring or remote mirroring takes advantage of the functionalities comprising:...

... In the identical manner, a source volume n, frozen at time t_{n-1} , will be supported by an auxiliary volume n, and an RVVn will exist.

Provisional Application, enablement support for Claim 1:

See page 6, line 9 to page 12, line 26:

Data Structure of a Freeze ...

...It is noted that instead of mirroring a source volume from a local SANL to a remote SANR, it is also feasible to mirror

a source volume from a local SANL to the very same SANL,
using the same mirroring procedure.

International Application, written description for Claim 1:

See page 1 line 5 - 7:

The invention relates in general to the field of mirroring,
or data replication, and in particular, to the asynchronous
mirroring of data objects between storage devices coupled to
a Storage Area Network (SAN) or to a network connectivity in
general.

See page 2, lines 24 - 6:

It is an object of the present invention to provide a method
and a system operative for mirroring a selected data object
from at least one local storage device (SDL) into at least
one remote storage device (SDRx). The at least one local
storage device is coupled to a first processing facility
(HL), and the at least one remote storage device is coupled
to a second processing facility (HR). The at least one local
storage device, the at least one remote storage device, the
first and the second processing facility are coupled to a
network connectivity comprising pluralities of users, of
processing facilities and of storage devices. The method and
the system comprise:

running a mirroring functionality in the first and in the
second processing facility, the mirroring functionality
comprising:

a freeze procedure for freezing the selected data
object,

a copy procedure for copying the frozen selected data
object into the at least one remote storage device,

permitting use and updating of the selected data object in
parallel to running the mirroring functionality, and

commanding, by default, repeated run of the mirroring
functionality for copying updates to the selected data
object, unless receiving command for mirroring break,
whereby the selected data object residing in the at least one
local storage device is copied and sequentially updated into
the at least one remote storage device.

See page 6 line 5 - 9:

The present invention achieves mirroring, or replication, of a selected data object from a local storage device, to a remote storage device, by sequential freeze and copy of discrete blocks of data. During mirroring, the selected data object may be used uninterruptedly, since mirroring is transparent to the operating system. Copying of the successive discrete blocks of data is performed asynchronously and in the background.

See page 6, line 20 to page 7 line 23:

When the mirroring functionality is operated, a selected data object is frozen by a freeze procedure, for example as a source volume ...

...Mirroring Overview Table.

International Application, enablement support for Claim 1:

See page 8, line 16 to page 20 line 26:

The Freeze Procedure...

...As a simple example for the method of row I, both the local host HL and the remote host HR operate the mirroring functionality, each host acting as the local host while the other host is the remote host.

Claim 2

Provisional Application, written description support for

Claim 2:

See page 2, line 21 - 24:

The aim of the mirroring functionality is to create the replication of an object. For example, mirroring an object from one storage device to another storage device in the same SAN or remote mirroring of an object residing in a local site SAN (SANL) to any remote site SAN (SANRi), or vice-versa.

See page 3, line 1 - 5:

In addition to remote mirroring one-to-one, thus from one local SAN to one remote SAN, a one-to-many remote mirroring capability is also available. The remote mirroring operation from one-to-many, thus from one local SAN to many remote SANs, applies the method for remote one-to-one mirroring many times, for each remote mirroring operation from the local SAN to each remote SAN.

See page 4, line 1 to page 5 line 19:

Mirroring or remote mirroring takes advantage of the functionalities comprising: ...
...Write commands are forwarded to and carried out solely to the auxiliary volume.

Provisional Application, enablement support for Claim 2:

See page 4, lines 16 - 27:

Upon initiation of a freeze by of the DMVA (Data Mover and Virtualization Appliance), the VA (Virtualization Agent) takes a "freeze" of, for example, an object such as volume, a data file, a Data Base record or even a VV (Virtual Volume). A freeze is thus a copy of the contents of say, a VV taken at a given instant in time. Such a freeze previously taken (copy), thus permits the reconstruction of a lost VV with exactly the same contents that was frozen in time at the moment the freeze was taken.

See page 4, line 28 to page 5, line 19:

When a DMVA requests a freeze, for example, of an existing volume, the VA creates a RVV (Resulting Virtual Volume) consisting of two volumes; the first volume being the freeze of the existing volume now called Source Volume, and a newly created Virtual Volume labeled as the Auxiliary Volume. The source volume is thus the actual volume from which the freeze is taken and the auxiliary volume is a VV that is required for storing of newly created data as well as of a mapping table.

At the moment of the freeze, at time t, the source volume is frozen in the sense that any further operation that may alter the contents of the source volume is directed away therefrom. Read operations are thus allowed to be executed, as opposed to Write operations, which are prevented since they effect changes to the contents. It should be noted that the performance of the SAN is

only but slightly affected by the freeze functionality, and that from the O.S. point of view, the RVV as a whole, is seen as one single source volume. It is kept in mind that there is no controller on the data path. The activity required is restricted to the redirection of copy or write instructions.

During the operation of the SAN, from the moment the object freeze procedure starts, the freeze software module that resides in the VA, intercepts every I/O command addressed to that object and redirects that I/O command either to the source volume, or to the auxiliary volume, as will be explained below. From the instant that the RVV is created, the source volume is restricted to "Read Only" commands, in as far as the data to be read was not amended after the beginning of the freeze procedure. Write commands are forwarded to and carried out solely to the auxiliary volume.

See page 10, line 21 to page 12 line 23:

When mirroring is ordered, as in step 202, for example for a certain volume FL residing at the local site SANL, to be mirrored to a given remote site SRx, then the next step, here 204, is to set a cascade counters to one ($s = 1$)...

...Finally, the last stage E, illustrated in Fig. 8, is equivalent to an additional loop, after passing again through step 222. Here, at the local SAN SANL, the VV AVL3 is running and the frozen VV AVL2 is ready to be copied to the just-created and equally-sized VV AVR_x/2. Since the control loop repeats the run from step 218 to step 236 in Fig. 7, the repetitive stepwise procedure is called cascading.

International Application, written description for Claim 2:

See page 3, lines 3 - 15:

It is a further object of the present invention to provide a method and a system for

- applying the freeze procedure for freezing the selected data object as a source volume (SV),

- creating at least one local auxiliary volume (AVL) to which updates addressed to the selected data object are redirected, each single data object out of the selected data object corresponding to one volume out of the at least one auxiliary volume,

- creating at least one remote volume in each remote storage device out of the at least one remote storage device, to correspond to each one local auxiliary volume created,

forming in the at least one local storage device of at least one resulting source volume, comprising the frozen selected data object and the at least one local auxiliary volume, and

applying the copy procedure for copying the frozen selected data object from the at least one resulting volume into the at least one remote storage device.

See page 6, line 20 to page 7 line 4:

When the mirroring functionality is operated, a selected data object is frozen by a freeze procedure, for example as a source volume ...

...A singular case relates to the mirroring of a selected data object consisting of only a single data object, residing in one local storage device, to but one remote storage device.

International Application, enablement support for Claim 2:

See page 8, line 17 to page 20 line 26:

A graphical illustration of the freeze procedure is depicted in Fig. 2, in stages from 2a to 2d. The horizontal axis t refers to time, starting with t=0...

...As a simple example for the method of row I, both the local host HL and the remote host HR operate the mirroring functionality, each host acting as the local host while the other host is the remote host.

Claim 3

Provisional Application, written description support for

Claim 3:

See page 3, lines 1 - 5:

In addition to remote mirroring one-to-one, thus from one local SAN to one remote SAN, a one-to-many remote mirroring capability is also available. The remote mirroring operation from one-to-many, thus from one local SAN to many remote SANs, applies the method for remote one-to-one mirroring many times, for each remote mirroring operation from the local SAN to each remote SAN.

Provisional Application, enablement support for Claim 3:

See page 10, line 14 to page 11, line 32:

The cascading functionality is now explained below, referring to the successive application of the virtualization of the freeze functionalities, thus generating a repetitive, cascade-like implementations of these functionalities...

...The procedure repeats a loop through the steps from 220 to 234 inclusive, and either continues to cascade or else, stops if so commanded.

See pages 12, line 1 - 23:

Another illustration of the mirroring application is depicted in Fig. 8. A separate column is reserved for the activities at the local site SANL on the left side, and at the remote site SANRx, on the right side, with five successive stages designated as A to E inclusive. Stage A relates to the source volume FL at the SANL, prior to the remote mirroring command given at step 204 of Fig. 7. At this stage, the cascade counter s is zero and the volume FL is operative. At stage B, with the cascade counter at s=1, the VV AVL1 is already created and running and the VV FM, at the remote site has just been created. Hence, this condition corresponds to the end of step 208 in Fig. 7. It is noted that the size of the VV FM is the same as that of the object source volume FL.

The picture shown at stage C, for s=2, is comparable to the end of step 222 in Fig. 7. At the local site SANL, the VV AVL1 is frozen, the VV AVL2 has been created and is now operative. In parallel, at the remote site SANRx, the VV AVRx/1 has just been created, with a size equal to that of the source volume FL.

One more development is depicted in stage D of Fig. 8, still with the cascade counter at s=2, in association with the end of step 234 of Fig. 7. The VV AVL1 is already resynchronized with the source volume FL and likewise, the VV AVRx/1 is resynchronized with the VV FM.

Finally, the last stage E, illustrated in Fig. 8, is equivalent to an additional loop, after passing again through step 222. Here, at the local SAN SANL, the VV AVL3 is running and the frozen VV AVL2 is ready to be copied to the just-created and equally-sized VV AVRx/2. Since the control loop repeats the run from step 218 to step 236 in Fig. 7, the repetitive stepwise procedure is called cascading."

International Application, written description for Claim 3:

See page 3, lines 16 - 18:

The mirroring functionality is applied simultaneously to more than one data object, and from at least one local storage device to at least one remote storage device, and vice-versa.

See page 6, line 38 to page 7 line 4:

Likewise, the data object may be used with the previous resulting source volume to which the last frozen local auxiliary volume is added to form a last resulting source volume. In principle, the mirroring functionality performs successive freeze and copy procedures to replicate one, or a group of data object(s), from one or more local storage device(s), to one or more other, or remote, storage device(s). A singular case relates to the mirroring of a selected data object consisting of only a single data object, residing in one local storage device, to but one remote storage device.

See page 7 lines 5 - 9:

The mirroring functionality is operable to perform more than one mirroring operation simultaneously. For example, two different data objects, each one residing in say, a different volume in a different local storage device, are possibly mirrored to two different remote storage devices. Evidently, simultaneous mirroring is not limited to two selected data objects.

See page 19, line 5 to page 20 line 26:

This is the simplest and basic mirroring method implementation for mirroring one data object, from one local storage device to one remote storage device...

...As a simple example for the method of row I, both the local host HL and the remote host HR operate the mirroring functionality, each host acting as the local host while the other host is the remote host.

International Application, enablement support for Claim 3:

See page 8, lines 7 - 10:

Mirroring of a selected data object residing in more than one storage devices may be effected to one or more storage devices. The minimum requirements are for two processing

facilities and for at least two storage devices on the network connectivity: one local storage device for copying from and one remote storage device for writing thereto.

See page 8, line 19 - 26:

It is assumed that the mirroring functionality operates on at least two processing facilities, such as a first and a second processing facility, respectively HL and HR, coupled to a network connectivity NET. A first storage device SDL and at least one second storage device SDR_x, where x identifies the specific storage device, referred to as, respectively, the local storage device and the at least one remote storage device, as are also coupled to the network connectivity NET. The at least one remote storage device SDR_x may thus consist of a first remote storage device SDR₁, a second remote storage device (SDR₂) and so on.

Claim 4

Provisional Application, written description support for Claim 4:

See page 2, line 21 - 31:

The aim of the mirroring functionality is to create the replication of an object. For example, mirroring an object from one storage device to another storage device in the same SAN or remote mirroring of an object residing in a local site SAN (SANL) to any remote site SAN (SANR_i), or vice-versa. The reverse operation, thus mirroring from a remote site SAN SANR_i to a local site SAN SANL, is achieved in exactly the same fashion by mere switching of the name designation of the two SAN sites involved. Thus, to replicate from a SANR_i to the SANL, rename the SANR_i as local, to SANL and the SANL as remote, to SANr_i. This cross-replication capability is not restricted to one specific local and one other precise remote site but is extendable to any couple of sites selected as a local site and a remote site, out of the sites of the MSAN.

Provisional Application, enablement support for Claim 4:

See page 10, line 24 to page 12, line 23:

When mirroring is ordered, as in step 202, for example for a certain volume FL residing at the local site SANL, to be mirrored to a given remote site SR_x, then the next step, here 204, is to set a cascade counters to one (s = 1)...

... Since the control loop repeats the run from step 218 to step 236 in Fig. 7, the repetitive stepwise procedure is called cascading.

International Application, written description for Claim 4:

See page 3, lines 16 - 18:

The mirroring functionality is applied simultaneously to more than one data object, and from at least one local storage device to at least one remote storage device, and vice-versa.

See page 7, lines 5 - 9:

The mirroring functionality is operable to perform more than one mirroring operation simultaneously. For example, two different data objects, each one residing in say, a different volume in a different local storage device, are possibly mirrored to two different remote storage devices. Evidently, simultaneous mirroring is not limited to two selected data objects."

International Application, enablement support for Claim 4:

See page 7, lines 31 to page 8, line 12:

Under the control of a first, or local processing facility, mirroring may take place from one local storage device to another remote storage device controlled by a second, or remote processing facility...

...As stated above, the mirroring of a data object from one storage device to another storage device requires the application of successive freeze and copy procedures."

See page 8, line 19, to page 9, line 7:

It is assumed that the mirroring functionality operates on at least two processing facilities, such as a first and a second processing facility, respectively HL and HR, coupled to a network connectivity NET...

...From the point of view of the Operating System O.S., the resulting source volume is seen as the original selected data object with which is used transparently.

See page 9, line 24, to page 10, line 19:

Referring to stage 2b of Fig. 2, at time $t = 2$, after the frozen source volume SV is copied, a next freeze is performed and applied to the local auxiliary volume 1 AVL1...

...Evidently, use of the data object is permitted to continue, in association with the ultimate resulting source volume consisting of the last resulting source volume and of the ultimate local auxiliary volume.

See page 14, line 11, to page 20, line 26:

Referring to the description related to the freezing of a source volume SV, at stage 2a in Fig. 2, it was stated that the source volume was copied after the freeze took place...

... As a simple example for the method of row I, both the local host HL and the remote host HR operate the mirroring functionality, each host acting as the local host while the other host is the remote host.

Claim 5

Provisional Application, written description support for Claim 5:

Not certain

Provisional Application, enablement support for Claim 5:

Not certain.

International Application, written description for Claim 5:

See page 3, lines 19 - 21:

It is another object of the present invention to provide a method and a system for:
applying the freeze procedure for freezing simultaneously more than one data object,...

See page 6, line 5 to page 7, line 23:

The present invention achieves mirroring, or replication, of a selected data object from a local storage device, to a remote storage device, by sequential freeze and copy of discrete blocks of data...

...Mirroring Overview Table.

International Application, enablement support for Claim 5:

See page 6, line 5 to page 7, line 23:

The present invention achieves mirroring, or replication, of a selected data object from a local storage device, to a remote storage device, by sequential freeze and copy of discrete blocks of data ...

...Mirroring Overview Table.

Claim 6

Provisional Application, written description support for Claim 6:

See page 2, line 21 to page 3 line 5:

The aim of the mirroring functionality is to create the replication of an object ...

...The remote mirroring operation from one-to-many, thus from one local SAN to many remote SANs, applies the method for remote one-to-one mirroring many times, for each remote mirroring operation from the local SAN to each remote SAN.

Provisional Application, enablement support for Claim 6:

See page 10, line 21 to page 12 line 26:

Fig. 7 illustrates the consecutive steps of the mirroring functionality referring to the local SAN, SANL, in the left column and to the remote SANRx in the right column...

...It is noted that instead of mirroring a source volume from a local SANL to a remote SANR, it is also feasible to mirror a

source volume from a local SANL to the very same SANL, using the same mirroring procedure.

International Application, written description for Claim 6:

See page 3, lines 22 - 23:

Applying the copy procedure to copy simultaneously more than one frozen selected data object, ...

See page 6, line 32 to page 7 line 31:

Once the frozen source volume is mirrored to the remote storage device, the freeze and copy procedures are repeated.

...Mirroring Overview Table.

See page 8, lines 27 - 33:

The designations local and remote are used for origin and destination, without implying any restriction on the physical location of the storage devices. Thus, both the local and the remote storage devices may reside, say, inside the same or in different storage device(s) coupled to a SAN, or to a host H, the different storage devices being adjacent or each one on opposite side of the globe. Copy is made from the local storage device to one or more remote storage device(s). Any storage device may be designated with either name, but there is only one local storage device when mirroring therefrom.

International Application, enablement support for Claim 6:

See page 14, line 25 to page 17 line 2:

An illustration of the mechanisms of the mirroring functionality is presented in Fig. 7 as a general overview, while a more detailed description is provided with reference to Fig. 8.

In Fig. 7 the left column relates to the local storage device SDL wherein a data object resides in the source volume SV, and the abscise displays a time axis t. The right column indicates events occurring in parallel to those at the local storage device, and depicts the process at the remote storage device SDR_x, where x[1, 2, ..., n] is chosen out of the at least one x available storage device. The denomination "the

remote storage device SDRx" is used below in the sense of at least one storage device.

Stage 7A in Fig. 7 shows the situation prior to mirroring. In the left column, the source volume SV created at time $t = 0$ contains the data object, while a mirroring cycle counter s is at zero. There are no events in the right column.

At stage 7B, in the left column, the mirroring counter is increased by one to $s = 1$ and a freeze of the source volume SV is commanded at time $t = 1$. At the same time, a first local auxiliary volume 1 AVL1 is created in the local storage device SDL, whereto updates to the data object are now directed. The updates are those I/O WRITE instructions from the computing facility HL that are redirected to the local auxiliary volume.

Simultaneously with the freeze at $t = 1$, a first remote volume RVx/s, here RVx/1, is created in the remote storage device SDRx, in the right column of Fig. 7, with the same size as the source volume SV. In turn, the frozen source volume SV is copied, in the background, and written to the remote volume RVx/1.

It was stated above that the freeze procedure divides a frozen data object into chunks of e.g. 1 MB. Upon creation of a local auxiliary volume and of the resulting source volume, a freeze table is also created therein, to relate between the source volume and the updates. The freeze table redirects I/O instructions from the data object to the local auxiliary volume, when necessary.

Meanwhile, the O.S. remains in operative association with both the source volume SV and the first local auxiliary volume AVL1, forming together the resulting source volume. It is noted that mirroring is executed in the background without need to wait for I/O instructions from the remote storage device. Thereby, the speed of operation of the local processor HL or of the network facility is not impaired.

At stage 7C of Fig. 7, the mirroring counter is increased by one to $s = 2$ and a second freeze command is received at time $t = 2$, occurring at or after completion of the copy operation of the source volume SV to the first remote volume RVx/1. Simultaneously, the first local auxiliary volume AVL1 is frozen and a second remote volume RVx/2 is created in the remote storage device SDRx, in the right column, with the same size as the first local auxiliary volume AVL1. A second

local auxiliary volume 2 AVL2 is created in the local storage device SDL whereto updates to the data object are directed.

A freeze table is automatically created by the freeze procedure, to reside in each local auxiliary volume, to the advantage of the O.S. In turn, the first local auxiliary volume AVL1, including the freeze tables for the benefit of the second computing facility HR, is copied to and written to the second remote volume Rvx/2.

At the same time, a new resulting source volume is created together with a new freeze table. The new resulting source volume consists of the previous resulting source volume to which is added the second local auxiliary volume AVL2. The O.S may thus communicate with the new resulting source volume to use the data object in parallel to mirroring.

At time $t = 2$ in the left column of stage 7C, the local storage device SDL contains the source volume SV, the first local auxiliary volume AVL1 and the second local auxiliary volume AVL2. At the same time in the right column, the remote storage device SDRx contains the first and the second remote volumes.

Still with the mirroring counter at $s = 2$, but at stage 7D, the frozen volumes, namely the source volume SV and the first local auxiliary volume are synchronized, whereby the updates previously written into the first local auxiliary volume AVL1 are entered into the source volume SV. The freeze table residing in the first local auxiliary volume AVL1 is used for correctly synchronizing the updates. The first local auxiliary volume AVL1, which contains at most as many chunks or segments as the source volume SV, is copied to overwrite the contents of the source volume SV that retains its original size. The first local auxiliary volume AVL1 is now deleted.

The indices opposite the chunk numbers in the freeze table residing in the second local auxiliary volume AVL2 are set to index values of -1, to reflect the status of the synchronized volumes. In parallel, the second remote volume RVx/2 is synchronized into the first volume RVx/1, which retains the same size as the source volume SV. Synchronization at the remote storage device is performed by the second processing facility HR using the freeze table copied thereto together with the last copied local auxiliary volume. The second remote volume RVx/2 may now be deleted.

Synchronization limits the required storage space in both the local storage device SDL and the remote storage device SDR_x, by deleting the local auxiliary volume and the remote volume that now becomes unnecessary.

Stage 7E is another freeze stage, equivalent to stages 7B and 7C. The mirroring cycle counter at the first computing facility HL is increased by one to $s = 3$, and a freeze of the second local auxiliary volume AVL2 is executed at time $t = 3$. In addition, a third local auxiliary volume AVL3 is created with the local storage device SDL. Simultaneously, a third remote volume RV_x/3 is created in the remote storage device SDR_x, in the right column, with the same size as the second auxiliary volume AVL2. The ultimate resulting source volume now contains the previous resulting source volume plus the ultimate local auxiliary volume AVL3.

As before, the last frozen local auxiliary volume, here AVL2, is copied to the last created remote volume, RV_x/3. After copy completion is acknowledged to the first computing facility HL, command is given to synchronize the last frozen local auxiliary volume AVL2 with the source volume SV.

In the remote storage device SDR_x, the second remote volume RV_x/2 is synchronized with the first remote volume, RV_x/1, under control of the second computing facility HR. It is noted that at this third mirroring cycle for $s = 3$, the remote storage device SDR_x now contains a copy of the resulting source volume that existed in the first mirroring cycle, at $s = 1$. At a mirroring cycle of $s = T$, the copy saved in the remote storage device SDR_x is always that of the resulting source volume at time $s = T-2$. At all times, there is a lag of two mirroring cycles between the last held copy at the remote storage device and the ultimate resulting source volume in the local storage device SDL.

Next, the process continues in the same manner as described above.

It is noted that the denomination remote storage device x , SDR_x, is a name used to refer to a storage device different from the local storage device, at the same site or at a remote site. Thus, mirroring from a source volume SV residing in a local SANL at a local site, is feasible not only to a storage device at the local site, but also to a storage device emplaced at a remote site, using the same mirroring procedure. Likewise, cross mirroring is feasible, as well as simultaneous cross mirroring.

See page 19, line 41 to page 20 line 26:

Row III of Table 2 calls for the mirroring of a selected data object residing in local storage SDL as single data objects, thus as a group of data objects, into one remote storage device SDRx. The mirroring functionality is applied as by row I, by freezing all the single data objects simultaneously. For example, if the selected data object is a group of three single data objects, then these three are frozen at the same time, and then each one is copied to the remote storage device SDRx. The next mirroring cycle may now start after completion of writing to the storage device SDRx.

Row IV presents the issue of mirroring a selected data object consisting of e.g. three single data objects residing in three different local volumes SDL_i, with $i = 1, 2, \text{ and } 3$, to one remote storage device SDRx. Again, the freeze procedure is simultaneous for the three single data objects and the method of row I is applied to each one of the three single data objects. A next mirroring cycle will start after completion of the last write operation to the destination remote storage device SDRx.

Row V applies the freeze procedure as by the method of row III and the copy procedure for copy to many remote storage devices as by row II.

An example for the mirroring of a selected data object consisting of a group of single data objects residing in a group of storage devices, with the number of single data objects being equal to the number of destination remote storage devices, is shown in Row IV. The simultaneous freeze for more than one data object is similar to the freeze procedure applied in row III, and the copy procedure is similar the one applied in row II.

It is important to note that the freeze procedure is simultaneous for all more than one data objects to be frozen, whether belonging to the same selected data object or stored in more than one local storage device. The cycle time to the next mirroring cycle is dictated by the time needed for the copy procedure to complete the last copy, when multiple copies are performed, such as to many remote storage devices.

It is also noted that simultaneous cross mirroring, from local to remote storage device and vice-versa is also practical with the mirroring functionality for the rows I to VI inclusive. As a simple example for the method of row I, both the local host HL and the remote host HR operate the

mirroring functionality, each host acting as the local host while the other host is the remote host.

Claim 7

Provisional Application, written description support for Claim 7:

See page 3 lines 1 - 5:

In addition to remote mirroring one-to-one, thus from one local SAN to one remote SAN, a one-to-many remote mirroring capability is also available. The remote mirroring operation from one-to-many, thus from one local SAN to many remote SANs, applies the method for remote one-to-one mirroring many times, for each remote mirroring operation from the local SAN to each remote SAN.

Provisional Application, enablement support for Claim 7:

See page 3 lines 6 - 8:

The description of the implementation of one-to-one remote mirroring is presented below with reference to a local site SANL and to a specific remote site SANRx...

See page 10, line 14 to page 12 line 26:

The cascading functionality is now explained below, referring to the successive application of the virtualization of the freeze functionalities, thus generating a repetitive, cascade-like implementations of these functionalities...

...It is noted that instead of mirroring a source volume from a local SANL to a remote SANR, it is also feasible to mirror a source volume from a local SANL to the very same SANL, using the same mirroring procedure.

International Application, written description for Claim 7:

See page 3, lines 19 - 31:

It is another object of the present invention to provide a method and a system for:
 applying the freeze procedure for freezing simultaneously more than one data object,
 applying the copy procedure to copy simultaneously more than one frozen selected data object,

mirroring simultaneously one single data object residing in one local storage device into more than one remote storage device,

mirroring simultaneously a plurality of single data objects residing respectively in a same plurality of local storage devices into one remote storage device,

mirroring simultaneously a plurality of single data objects residing in one local storage device respectively into a same plurality of remote storage devices, and

mirroring simultaneously one single data object residing in each one local storage device out of a plurality of local storage devices into one remote storage device.

See page 6, lines 24 - 25:

mirroring simultaneously one single data object residing in one local storage device into more than one remote storage device,...

See page 6, line 40 to page 41 line 23:

In principle, the mirroring functionality performs successive freeze and copy procedures to replicate one, or a group of data object(s), from one or more local storage device(s), to one or more other, or remote, storage device(s)...

... Mirroring Overview Table

International Application, enablement support for Claim 7:

See page 7, line 31, to page 8 line 12:

Under the control of a first, or local processing facility, mirroring may take place from one local storage device to another remote storage device controlled by a second, or remote processing facility...

... As stated above, the mirroring of a data object from one storage device to another storage device requires the application of successive freeze and copy procedures.

Claim 8

Provisional Application, written description support for Claim 8:

See page 3, lines 1 - 5:

In addition to remote mirroring one-to-one, thus from one local SAN to one remote SAN, a one-to-many remote mirroring capability is also available. The remote mirroring operation from one-to-many, thus from one local SAN to many remote SANs, applies the method for remote one-to-one mirroring many times, for each remote mirroring operation from the local SAN to each remote SAN.

Provisional Application, enablement support for Claim 8:

See page 5, line 20 to page 11, line 28:

A graphical illustration of the freeze facility is depicted in Fig 2 where the progress of a series of consecutive freeze procedures is shown. The horizontal axis refers to time, running from $t-1$, t_0 , t_1 , t_2 , to t_n

... Such resynchronization and unification is performed, respectively, in steps 232 and 234, whereby only FL and FM remain available, while both AVL1 and AVR_x/1 are deleted. If so wished, the mirroring loops may be broken in step 236 and ended in step 238, or else, continued.

International Application, written description for Claim 8:

See page 3, lines 16 - 18:

The mirroring functionality is applied simultaneously to more than one data object, and from at least one local storage device to at least one remote storage device, and vice-versa.

See page 8, lines 7 - 8:

Mirroring of a selected data object residing in more than one storage devices may be effected to one or more storage devices.

See page 17, lines 1 - 2:

Likewise, cross mirroring is feasible, as well as simultaneous cross mirroring."

See page 19, Table 2, III to VI:

	MIRROR	FROM		TO	
	Data Objects	Local Storage Devices	Created Local Auxiliary Volumes Per mirroring Cycle =	Remote Storage Devices	Maximum # of Created Remote Volumes Per mirroring Cycle =
I	1	1	1	1	1
II	1	1	1	>1	# of Remote Storage Devices
III	>1	1	Data Objects	1	# of Data Objects
IV	>1	>1	Data of Objects or # of Local Storage Devices	1	# Data of Objects or # of Local Storage Devices
V	>1	1	Data Objects	>1	# of Data Objects or # of Remote Storage Devices
VI	>1	>1	Data Objects or # of Local Storage Devices	>1	# of Data Objects or # of Local Storage Devices

International Application, enablement support for Claim 8:

See page 17, lines 16 - 19:

Mirroring is directed to one or more storage devices referred to as remote storage device x, SDR_x, where x is an integer, from 1 to n. Control passes first to step 204, where a mirroring cycle counter s is set to s = 1, and continues to step 206.

See page 19, line 8 to page 20, line 26:

For example, in row II, one data object is stored in one local storage device SDL, for mirroring into a plurality of remote storage devices SDR_x, where x receives the identity of the specific storage device, will require the creation of a number of remote volumes equal to the number of the plurality of remote storage devices, for each mirroring cycle...

...As a simple example for the method of row I, both the local host HL and the remote host HR operate the mirroring functionality, each host acting as the local host while the other host is the remote host.

Claim 9

Provisional Application, written description support for Claim 9:

See page 3 lines 1 - 5:

In addition to remote mirroring one-to-one, thus from one local SAN to one remote SAN, a one-to-many remote mirroring capability is also available. The remote mirroring operation from one-to-many, thus from one local SAN to many remote SANs, applies the method for remote one-to-one mirroring many times, for each remote mirroring operation from the local SAN to each remote SAN.

Provisional Application, enablement support for Claim 9:

See page 10, line 21 to page 11, line 32:

Fig. 7 illustrates the consecutive steps of the mirroring functionality referring to the local SAN, SANL, in the left column and to the remote SANRx in the right column...

... The procedure repeats a loop through the steps from 220 to 234 inclusive, and either continues to cascade or else, stops if so commanded.

International Application, written description for Claim 9:

See page 3 lines 19 - 31:

It is another object of the present invention to provide a method and a system for:

- applying the freeze procedure for freezing simultaneously more than one data object,
- applying the copy procedure to copy simultaneously more than one frozen selected data object,
- mirroring simultaneously one single data object residing in one local storage device into more than one remote storage device,
- mirroring simultaneously a plurality of single data objects residing respectively in a same plurality of local storage devices into one remote storage device,

mirroring simultaneously a plurality of single data objects residing in one local storage device respectively into a same plurality of remote storage devices, and
mirroring simultaneously one single data object residing in each one local storage device out of a plurality of local storage devices into one remote storage device.

See page 7 lines 5 - 9:

The mirroring functionality is operable to perform more than one mirroring operation simultaneously. For example, two different data objects, each one residing in say, a different volume in a different local storage device, are possibly mirrored to two different remote storage devices. Evidently, simultaneous mirroring is not limited to two selected data objects.

See page 7, lines 21 - 23, item IV:

	MIRROR	FROM	TO
	# of Data Objects	Local Storage Devices	Remote Storage Devices
I	1	1	1
II	1	1	>1
III	>1	1	1
IV	>1	>1	1
V	>1	1	>1
VI	>1	>1	>1

International Application, enablement support for Claim 9:

See page 17, line 4 to page 20 line 26:

Fig. 8 illustrates the consecutive steps of the mirroring functionality, applicable to any network connectivity...

As a simple example for the method of row I, both the local host HL and the remote host HR operate the mirroring functionality, each host acting as the local host while the other host is the remote host.

Claim 10

Provisional Application, written description support for Claim 10:

Not certain.

Provisional Application, enablement support for Claim 10:

Not certain.

International Application, written description for Claim 10:

See page 3 lines 28 - 29:

Mirroring simultaneously a plurality of single data objects residing in one local storage device respectively into a same plurality of remote storage devices, and ...

See page 7 lines 5 - 9:

The mirroring functionality is operable to perform more than one mirroring operation simultaneously. For example, two different data objects, each one residing in say, a different volume in a different local storage device, are possibly mirrored to two different remote storage devices. Evidently, simultaneous mirroring is not limited to two selected data objects.

See page 7, lines 21.- 23, item VI:

	MIRROR	FROM	TO
	# of Data Objects	Local Storage Devices	Remote Storage Devices
I	1	1	1
II	1	1	>1
III	>1	1	1
IV	>1	>1	1
V	>1	1	>1
VI	>1	>1	>1

International Application, enablement support for Claim 10:

See page 17, line 4 to page 20 line 26:

Fig. 8 illustrates the consecutive steps of the mirroring functionality, applicable to any network connectivity...

...As a simple example for the method of row I, both the local host HL and the remote host HR operate the mirroring functionality, each host acting as the local host while the other host is the remote host.

Claim 11

Provisional Application, written description support for Claim 11: Probably none.

Provisional Application, enablement support for Claim 11:

Probably none.

International Application, written description for Claim 11:

See page 3 lines 16 - 31:

The mirroring functionality is applied simultaneously to more than one data object, and from at least one local storage device to at least one remote storage device, and vice-versa.

It is another object of the present invention to provide a method and a system for:

applying the freeze procedure for freezing simultaneously more than one data object,

applying the copy procedure to copy simultaneously more than one frozen selected data object,

mirroring simultaneously one single data object residing in one local storage device into more than one remote storage device,

mirroring simultaneously a plurality of single data objects residing respectively in a same plurality of local storage devices into one remote storage device,

mirroring simultaneously a plurality of single data objects residing in one local storage device respectively into a same plurality of remote storage devices, and

mirroring simultaneously one single data object residing in each one local storage device out of a plurality of local storage devices into one remote storage device.

See page 7 lines 5 - 9:

The mirroring functionality is operable to perform more than one mirroring operation simultaneously. For example, two different data objects, each one residing in say, a different volume in a different local storage device, are possibly mirrored to two different remote storage devices. Evidently, simultaneous mirroring is not limited to two selected data objects.

See page 20, lines 28 - 30:

For example, more combinations of selected data objects, local and remote storage devices may be considered.

International Application, enablement support for Claim 11:

See page 17, line 4 to page 20 line 26:

Fig. 8 illustrates the consecutive steps of the mirroring functionality, applicable to any network connectivity...

...As a simple example for the method of row I, both the local host HL and the remote host HR operate the mirroring functionality, each host acting as the local host while the other host is the remote host.

Claim 12

Provisional Application, written description support for Claim 12:

See page 5, line 3 to page 6 line 8:

At the moment of the freeze, at time t , the source volume is frozen in the sense that any further operation that may alter the contents of the source volume is directed away therefrom.

...In the identical manner, a source volume n , frozen at time t_{n-1} , will be supported by an auxiliary volume n , and an RVVn will exist.

Provisional Application, enablement support for Claim 12:

See page 10, line 27 to page 12 line 24:

When mirroring is ordered, as in step 202, for example for a certain volume FL residing at the local site SANL, to be mirrored to a given remote site SRx, then the next step, here 204, is to set a cascade counters to one ($s = 1$)...

...Since the control loop repeats the run from step 218 to step 236 in Fig. 7, the repetitive stepwise procedure is called cascading.

International Application, written description for Claim 12:

See page 3, line 32 to page 4 line 11:

It is yet another object of the present invention to provide a method and a system for:
at a selected point in time:
 starting a mirroring cycle,
 freezing the selected data object,
 creating at least one local auxiliary volume (AVL) in the at least one local storage device (SDL) and at least one remote volume (RV) in the at least one remote storage device (SDRx),
 forming at least one resulting source volume comprising the frozen selected data object and the local auxiliary volume (AVL), and
after the selected point in time:
 copying the frozen selected data object from the resulting source volume into the at least one remote volume until completion of copy,
 redirecting to the local auxiliary volume of the updates addressed to the selected data object,
 permitting use of the selected data object during mirroring, by associative operation with the resulting source volume, and
 repeating a next mirroring cycle by default command, after completion of copy to the at least one remote storage device, unless receiving command for mirroring break.

See page 6, line 20 to page 7 line 23:

When the mirroring functionality is operated, a selected data object is frozen by a freeze procedure, for example as a source volume...

...Mirroring Overview Table

International Application, enablement support for Claim 12:

See page 7, line 25 to page 20 line 26:

Reference is made to Fig. 1 of the co-pending patent application PCT/IL00/00309, entitled "Storage Virtualization in a Storage Network", by the same applicant, incorporated herewith by reference in whole, cited below as the '309 patent...

...As a simple example for the method of row I, both the local host HL and the remote host HR operate the mirroring

functionality, each host acting as the local host while the other host is the remote host.

Claim 13

Provisional Application, written description support for Claim 13:

See page 5, line 31 to page 6 line 8:

At time t_1 , a next freeze is performed and applied to the auxiliary volume 1, so that both the source volume 1 and the auxiliary volume 1 become a new source volume 2...

...In the identical manner, a source volume n , frozen at time t_{n-1} , will be supported by an auxiliary volume n , and an RVV n will exist.

See page 6, line 16 to page 10 line 5:

Since the configuration of the source volume 1 is intentionally frozen to prevent any change as from the time t_0 of the freeze, all the updates pertaining to that source volume 1, thus certainly all the I/O WRITE instructions directed thereto, are routed to the simultaneously created auxiliary volume 1...

...In step C7, at least one calculated index number is entered in the freeze table, into the first free chunk in the auxiliary volume.

Provisional Application, enablement support for Claim 13:

See page 10, line 24 to page 12 line 23:

When mirroring is ordered, as in step 202, for example for a certain volume FL residing at the local site SANL, to be mirrored to a given remote site SR x , then the next step, here 204, is to set a cascade counters to one ($s = 1$)...

...Here, at the local SAN SANL, the VV AVL3 is running and the frozen VV AVL2 is ready to be copied to the just-created and equally-sized VV AVR $x/2$. Since the control loop repeats the run from step 218 to step 236 in Fig. 7, the repetitive stepwise procedure is called cascading.

International Application, written description for Claim 13:

See page 4, lines 12 - 34:

It is yet an object of the present invention to provide a method and a system for:

- starting a next mirroring cycle at a next point in time occurring after completion of copy to the at least one remote storage device (SDR),

- freezing the resulting source volume,

- creating an ultimate local auxiliary volume in the local storage device and an ultimate remote volume in the at least one remote storage device,

- forming an ultimate resulting source volume comprising the penultimate resulting source volume and the ultimate local auxiliary volume (AVL), and

after the next point in time:

- copying the penultimate local auxiliary volume into the ultimate remote volume, and,

- redirecting to the ultimate local auxiliary volume of the updates addressed to the selected data object,

- permitting use of the selected data object during mirroring, by associative operation with the ultimate resulting source volume, and

after completion of copy into the ultimate remote volume:

- synchronizing the penultimate local auxiliary volume into the frozen selected data object,

- synchronizing the at least one ultimate remote volume into the penultimate remote volume by command of the second processing facility (HR), and

- repeating, by default command, of a next mirroring cycle after completion of copy to the at least one second storage device, unless receiving command for mirroring break.

See page 6, line 32 to page 7 line 23:

Once the frozen source volume is mirrored to the remote storage device, the freeze and copy procedures are repeated...

...Mirroring Overview Table

International Application, enablement support for Claim 13:

See page 9, line 24 to page 10 line 19:

When a freeze of a source volume SV is ordered at time $t = 1$, the now frozen source volume is arbitrarily divided into sequentially numbered segments or chunks of 1 MB for example,

and these chunks are listed in a Freeze Table 1 created at freeze time within the local auxiliary volume 1 AVL1...

... Evidently, use of the data object is permitted to continue, in association with the ultimate resulting source volume consisting of the last resulting source volume and of the ultimate local auxiliary volume.

See page 14, line 11 to page 20 line 26:

Referring to the description related to the freezing of a source volume SV, at stage 2a in Fig. 2, it was stated that the source volume was copied after the freeze took place.

... As a simple example for the method of row I, both the local host HL and the remote host HR operate the mirroring functionality, each host acting as the local host while the other host is the remote host.

Claim 14

Provisional Application, written description support for Claim 14:

See page 5, line 31 to page 5, line 8:

At time t_1 , a next freeze is performed and applied to the auxiliary volume 1, so that both the source volume 1 and the auxiliary volume 1 become a new source volume 2...

...In the identical manner, a source volume n , frozen at time t_{n-1} , will be supported by an auxiliary volume n , and an RVV n will exist.

Provisional Application, enablement support for Claim 14:

See page 11, line 8 to page 12 line 23:

In the manner, as explained above, the VV AVL1, for $s = s-1$, is frozen next, before it is copied to the SANRx. Accordingly, for $s=2$, step 220 freezes the AVL($s-1$), thus AVL1, and creates a new AVLS, here AVL2...

... Since the control loop repeats the run from step 218 to step 236 in Fig. 7, the repetitive stepwise procedure is called cascading.

International Application, written description for Claim 14:

See page 4, line 35 to page 5 line 13:

It is still an object of the present invention to provide a method and a system for:
selecting still another point in time occurring after completion of copy of the penultimate local auxiliary volume, freezing the resulting source volume,
creating an ultimate local auxiliary volume in the local storage device and an ultimate remote volume in the at least one remote storage device,
forming an ultimate resulting source volume comprising the penultimate resulting source volume and the ultimate local auxiliary volume, and
copying the penultimate local auxiliary volume into the at least one ultimate remote volume,
redirecting to the ultimate local auxiliary volume of updates addressed to the selected data object,
permitting use of the selected data object during mirroring in associative operation with the ultimate resulting source volume,
synchronizing the penultimate local auxiliary volume into the selected data object,
synchronizing the at least one ultimate remote volume into the penultimate remote volume, and
repeating a next mirroring cycle by default command after completion of copy to the at least one second storage device, unless receiving command for mirroring break.

See page 9, line 39 to page 10 line 19:

The third resulting source volume thus consists of the first source volume SV as frozen at time $t = 1$, of the frozen local auxiliary volumes 1 and 2 respectively AVL1 and AVL2, and of the ultimate local auxiliary volume 3 AVL3...

...Evidently, use of the data object is permitted to continue, in association with the ultimate resulting source volume consisting of the last resulting source volume and of the ultimate local auxiliary volume.

International Application, enablement support for Claim 14:

See page 16, line 12 to page 18 line 33:

Synchronization limits the required storage space in both the local storage device SDL and the remote storage device

SDRx, by deleting the local auxiliary volume and the remote volume that now becomes unnecessary...

... The procedure repeats a loop through the steps from 218 to 236 inclusive, which either continues mirroring or else, ends mirroring if so commanded.

Claim 15

Provisional Application, written description support for Claim 15:

See page 10, line 27 to page 11 line 32:

Further on, in step 206, the volume FL is frozen and an Auxiliary Virtual Volume (VV), here AVL1 for $s = 1$, is created at the SANL, while in parallel another VV FM, with the same size as that of the volume FL, is created at the SANRx, as seen in step 208...

... If the mirroring loop is not broken, then control returns to step 218, where the cascade counters is increased again by 1, to $s=3$. The procedure repeats a loop through the steps from 220 to 234 inclusive, and either continues to cascade or else, stops if so commanded.

Provisional Application, enablement support for Claim 15:

See page 10, line 27 to page 11 line 32:

Further on, in step 206, the volume FL is frozen and an Auxiliary Virtual Volume (VV), here AVL1 for $s = 1$, is created at the SANL, while in parallel another VV FM, with the same size as that of the volume FL, is created at the SANRx, as seen in step 208.

... If the mirroring loop is not broken, then control returns to step 218, where the cascade counters is increased again by 1, to $s=3$. The procedure repeats a loop through the steps from 220 to 234 inclusive, and either continues to cascade or else, stops if so commanded.

International Application, written description for Claim 15:

See page 5, line 14 - 18:

It is yet a further object of the present invention to provide a method and a system for:

storing in the at least one remote storage device of a complete mirrored copy of the selected data object comprising updates entered thereto at the time when copy of the before to penultimate local auxiliary volume was completed.

...A second local auxiliary volume 2 AVL2 is created in the local storage device SDL whereto updates to the data object are directed.

See page 15, lines 19 - 23:

At stage 7C of Fig. 7, the mirroring counter is increased by one to $s = 2$ and a second freeze command is received at time $t = 2$, occurring at or after completion of the copy operation of the source volume SV to the first remote volume RVx/1. Simultaneously, the first local auxiliary volume AVL1 is frozen and a second remote volume RVx/2 is created in the remote storage device SDRx, in the right column, with the same size as the first local auxiliary volume AVL1. A second local auxiliary volume 2 AVL2 is created in the local storage device SDL whereto updates to the data object are directed.

See page 16, lines 23 - 26:

As before, the last frozen local auxiliary volume, here AVL2, is copied to the last created remote volume, RVx/3. After copy completion is acknowledged to the first computing facility HL, command is given to synchronize the last frozen local auxiliary volume AVL2 with the source volume SV.

See page 16, lines 31 - 32:

At a mirroring cycle of $s = T$, the copy saved in the remote storage device SDRx is always that of the resulting source volume at time $s = T-2$.

International Application, enablement support for Claim 15:

See page 17, line 35 to page 18 line 5:

By step 214, complementary to step 212, the source volume SV is written to the first remote volume RVx/1, and when ended, completion is acknowledged to the computing facility HL, which then performs a completion check in step 216, similarly to step 210. As before, a negative response causes a loop-again through the completion check at step 216, while a positive answer passes command to step 218, where the

mirroring cycle counter is increased by one, to $s = s + 1$, here $s = 2$.

Control is now forwarded to step 220, to continue mirroring. In the local storage device SDL there is created first an ultimate local auxiliary volume designated as AVL/s, which for $s = 2$, is the second local auxiliary volume AVL2, and then, the penultimate local auxiliary volume AVL/s-1, here AVL1, is frozen. There is also created an ultimate resulting source volume, in the manner described above."

See page 18, lines 13 -20:

In step 226 command is given to copy the frozen penultimate, here the first, local auxiliary volume AVL/s-1 to the ultimate, here the second, remote volume RVx/s. Step 228 executes the write operation from the first local auxiliary volume AVL/s-1 to the second remote volume RVx/s, which upon write completion, is acknowledged to step 230.

It is noted that at this stage, both the source volume SV and the first local auxiliary volume AVL1 are acknowledged as being actually mirrored to the SDRx, in both the RVx/1 and the AVRx/2. Meanwhile, the second freeze is operating at the local host HL and the new updates are redirected to the local auxiliary virtual volume AVL2."

Claim 16

Provisional Application, written description support for Claim 16:

See page 1, lines 15 - 18:

The description relates to the capability to replicate the contents of a storage device, pertaining to a local SAN (Storage Area Network), in part or in whole, to a remote-SAN coupled to the local SAN replication or mirroring understood as copying, are applicable to any object, volume, file, data structure or the like.

See page 1, lines 15 - 18:

The description relates to the capability to replicate the contents of a storage device, pertaining to a local SAN (Storage Area Network), in part or in whole, to a remote-SAN coupled to the local SAN replication or mirroring understood

as copying, are applicable to any object, volume, file, data structure or the like.

Provisional Application, enablement support for Claim 16:

See page 4, lines 16 - 19:

Upon initiation of a freeze by of the DMVA (Data Mover and Virtualization Appliance), the VA (Virtualization Agent) takes a "freeze" of, for example, an object such as volume, a data file, a Data Base record or even a VV (Virtual Volume).

International Application, written description for Claim 16:

See page 1, lines 9 - 12:

A selected data object is a single data object, or a plurality, or a group, of data objects.

A data object is a volume, a logical or virtual volume, a data file, or any data structure. The terms data object and volume are used interchangeably below.

International Application, enablement support for Claim 16:

See page 1, lines 9-12:

A selected data object is a single data object, or a plurality, or a group, of data objects.

A data object is a volume, a logical or virtual volume, a data file, or any data structure. The terms data object and volume are used interchangeably below.

See page 8, lines 34 - 40:

To start, the mirroring functionality, which contains both the freeze procedure and the copy procedure, receives indication of the data object selected to be frozen. As illustrated at stage 2a of Fig. 2, at a given moment, at time $t = 1$, the freeze procedure receives a request to freeze a selected data object as a source volume SV. In consequence, the "frozen" source volume SV is thus restricted to "read only", which does not alter the contents of the source volume. The frozen source volume SV may now be copied as will be described below.

See page 17, line 13 to page 19 line 3:

In step 202 of Fig. 8, command is given to mirror a selected source volume SV, which resides in a local storage device SDL that is coupled to a local host HL...

... Table 2

Claim 17

Provisional Application, written description support for Claim 17:

See page 7, lines 34 - 37:

Reference is made to Fig. 1 of the co-pending patent application PCT/IL00/00309, entitled "Storage Virtualization in a Storage Network", by the same applicant, incorporated herewith by reference in whole, cited below as the '309 patent. Reference is also made to Fig. 1 in the present application, depicting a network connectivity NET.

See PCT/IL00/00309, page 12 line 40 to page 13 line 3:

The virtualization software operates on a system involving a SAN with a Network Switch and a Storage Virtualization Manager (SVM). By definition, a SAN includes an array of host computers coupled via a Network Switch to an array of storage devices, either by Fibre Channel (FC), or by TCP/IP, or by other links, to one or more storage devices. The SAN is linked to a user network, preferably an Ethernet, such as a LAN, a WAN, or the Internet, to which users are connected. A System Administrator, operating a workstation coupled to the user network, is a user privileged with access to the SVM, to exercises real-time management of the SAN.

Provisional Application, enablement support for Claim 17:

See page 4, lines 16 - 19:

Upon initiation of a freeze by of the DMVA (Data Mover and Virtualization Appliance), the VA (Virtualization Agent) takes a "freeze" of, for example, an object such as volume, a data file, a Data Base record or even a VV (Virtual Volume).

International Application, written description for Claim 17:

See page 1, lines 24 - 25:

A network connectivity is a Local Area Network (LAN), a Wide Area Network (WAN), or a Storage Area Network (SAN).

International Application, enablement support for Claim 17:

See page 8, lines 8 - 15:

The minimum requirements are for two processing facilities and for at least two storage devices on the network connectivity: one local storage device for copying from and one remote storage device for writing thereto.

As stated above, the mirroring of a data object from one storage device to another storage device requires the application of successive freeze and copy procedures. However, the operation of a network connectivity may not be hampered while mirroring. Therefore, the description below illustrates first the freeze procedure, then the operation of the system while the freeze procedure is running and last, the copy procedure.

See page 17, line 4 to page 19 line 3:

Fig. 8 illustrates the consecutive steps of the mirroring functionality, applicable to any network connectivity...

... Table 2

Claim 18

Provisional Application, written description support for Claim 18:

See page 3, lines 2 - 5:

The remote mirroring operation from one-to-many, thus from one local SAN to many remote SANs, applies the method for remote one-to-one mirroring many times, for each remote mirroring operation from the local SAN to each remote SAN.

Provisional Application, enablement support for Claim 18:

See page 10, line 17 to page 12 line 23:

Fig. 7 illustrates the consecutive steps of the mirroring functionality referring to the local SAN, SANL, in the left column and to the remote SANRx in the right column...

... Since the control loop repeats the run from step 218 to step 236 in Fig. 7, the repetitive stepwise procedure is called cascading.

International Application, written description for Claim 18:

See page 5 lines 18 - 26:

It is yet a further object of the present invention to provide a method and a system for:
repeating operation of the mirroring functionality at discrete repetition intervals of time defined as lasting at least as long as duration of copying of the ultimate local auxiliary volume to the ultimate remote volume,
synchronizing updates to overwrite the selected data object, and
synchronizing a later remote volume to overwrite the penultimate resulting first remote volume.

See page 6, lines 10 - 11:

Mirroring consists of a succession of freeze and copy procedures repeated sequentially in successive mirroring cycles.

See page 6, line 32 to page 7 line 2:

Once the frozen source volume is mirrored to the remote storage device, the freeze and copy procedures are repeated...

... In principle, the mirroring functionality performs successive freeze and copy procedures to replicate one, or a group of data object(s), from one or more local storage device(s), to one or more other, or remote, storage device(s).

See page 14, line 19 - 20:

The mirroring functionality may be repeated sequentially, or may be stopped after any freeze and copy cycle.

International Application, enablement support for Claim 18:

See page 17, line 4 to page 19 line 3:

Fig. 8 illustrates the consecutive steps of the mirroring functionality, applicable to any network connectivity...

... Table 2

Claim 19

Provisional Application, written description support for Claim 19:

See page 6, lines 4 - 7:

When so decided, at the given freeze time t2, the source volume 2 will resynchronize, or unite with the auxiliary volume 2 to form a new source volume 3 and a new auxiliary volume 3, not seen in Fig. 2, will come to existence, forming a RVV2.

Provisional Application, enablement support for Claim 19:

See page 11, lines 23 - 27:

Practically, there is no further reason to operate either AVL1 or AVRx/1 separately, and therefore, those VVs may be resynchronized with, respectively, FL and FM. Such resynchronization and unification is performed, respectively, in steps 232 and 234, whereby only FL and FM remain available, while both AVL1 and AVRx/1 are deleted."

International Application, written description for Claim 19:

See page 5, lines 19 - 26:

It is yet a further object of the present invention to provide a method and a system for:
repeating operation of the mirroring functionality at discrete repetition intervals of time defined as lasting at least as long as duration of copying of the ultimate local auxiliary volume to the ultimate remote volume,
synchronizing updates to overwrite the selected data object, and
synchronizing a later remote volume to overwrite the penultimate resulting first remote volume.

International Application, enablement support for Claim 19:

See page 9, line 39 to page 10 line 17:

The third resulting source volume thus consists of the first source volume SV as frozen at time $t = 1$, of the frozen local auxiliary volumes 1 and 2 respectively AVL1 and AVL2, and of the ultimate local auxiliary volume 3 AVL3...

... At the same time, the updates residing in the before-penultimate local auxiliary volume AVL, here the local auxiliary volume 2 AVL2, will be synchronized into the first frozen source volume SV. The local auxiliary volume 2 AVL2 may now be deleted.

See page 15, line 37 to page 16 line 35:

Still with the mirroring counter at $s = 2$, but at stage 7D, the frozen volumes, namely the source volume SV and the first local auxiliary volume are synchronized, whereby the updates previously written into the first local auxiliary volume AVL1 are entered into the source volume SV...

... Next, the process continues in the same manner as described above.

Claim 20

Provisional Application, written description support for Claim 20:

See page 1, lines 15 - 18:

The description relates to the capability to replicate the contents of a storage device, pertaining to a local SAN (Storage Area Network), in part or in whole, to a remote-SAN coupled to the local SAN replication or mirroring understood as copying, are applicable to any object, volume, file, data structure or the like.

Provisional Application, enablement support for Claim 20:

See page 1, lines 15 - 18:

The description relates to the capability to replicate the contents of a storage device, pertaining to a local SAN (Storage Area Network), in part or in whole, to a remote-SAN coupled to the local SAN replication or mirroring understood as copying, are applicable to any object, volume, file, data structure or the like.

International Application, written description for Claim 20:

See page 1, lines 17 - 18:

The storage space for a data object may span only a part, or the whole, or more than the whole space contents of a storage device.

International Application, enablement support for Claim 20:

The storage space for a data object may span only a part, or the whole, or more than the whole space contents of a storage device.

Claim 21

Provisional Application, written description support for Claim 21:

See page 4, line 14 to page 5 line 19:

A freeze of an object, taken at time t , is a "frozen" copy of the contents of that object at that same time t ...

... Write commands are forwarded to and carried out solely to the auxiliary volume.

Provisional Application, enablement support for Claim 21:

See page 5, line 23 to page 6 line 8:

At a given moment t_0 in time, a request is made for the freeze of a selected object...

... In the identical manner, a source volume n , frozen at time t_{n-1} , will be supported by an auxiliary volume n , and an RVV n will exist.

International Application, written description for Claim 21:

See page 6, line 26 to page 7 line 4:

The selected data object may be used during mirroring. At freeze time, the Operating System O.S. creates a resulting source volume comprising both the frozen selected data object and the first local auxiliary volume...

...A singular case relates to the mirroring of a selected data object consisting of only a single data object, residing in one local storage device, to but one remote storage device.

International Application, enablement support for Claim 21:

See page 8, line 17 to page 18 line 33:

A graphical illustration of the freeze procedure is depicted in Fig. 2, in stages from 2a to 2d. The horizontal axis t refers to time, starting with $t = 0$...

...The procedure repeats a loop through the steps from 218 to 236 inclusive, which either continues mirroring or else, ends mirroring if so commanded.

Claim 22

Provisional Application, written description support for Claim 22:

See page 6, line 14 to page 6 line 8:

At time t_1 , a next freeze is performed and applied to the auxiliary volume 1, so that both the source volume 1 and the auxiliary volume 1 become a new source volume 2...

...In the identical manner, a source volume n, frozen at time t_{n-1} , will be supported by an auxiliary volume n, and an RRVn will exist.

See page 10, line 24 to page 12 line 23:

In the identical manner, a source volume n, frozen at time t_{n-1} , will be supported by an auxiliary volume n, and an RVVn will exist...

...Since the control loop repeats the run from step 218 to step 236 in Fig. 7, the repetitive stepwise procedure is called cascading.

Provisional Application, enablement support for Claim 22:

See page 10, line 24 to page 12 line 23:

In the identical manner, a source volume n, frozen at time t_{n-1} , will be supported by an auxiliary volume n, and an RVVn will exist...

... Since the control loop repeats the run from step 218 to step 236 in Fig. 7, the repetitive stepwise procedure is called cascading.

International Application, written description for Claim 22:

See page 4, lines 12 to page 5 line 34:

It is yet an object of the present invention to provide a method and a system for: ...

...repeating a next mirroring cycle by default command after completion of copy to the at least one second storage device, unless receiving command for mirroring break.

International Application, enablement support for Claim 22:

See page 9, line 24 to page 10 line 19:

Referring to stage 2b of Fig. 2, at time $t = 2$, after the frozen source volume SV is copied, a next freeze is performed and applied to the local auxiliary volume 1 AVL1...

...Evidently, use of the data object is permitted to continue, in association with the ultimate resulting source volume consisting of the last resulting source volume and of the ultimate local auxiliary volume.

See page 14, line 37 to page 16 line 14:

At stage 7B, in the left column, the mirroring counter is increased by one to $s = 1$ and a freeze of the source volume SV is commanded at time $t = 1$...

...Synchronization limits the required storage space in both the local storage device SDL and the remote storage device SDRx, by deleting the local auxiliary volume and the remote volume that now becomes unnecessary."

See page 17, line 35 to page 18 line 33:

By step 214, complementary to step 212, the source volume SV is written to the first remote volume RVx/1, and when ended, completion is acknowledged to the computing facility HL, which then performs a completion check in step 216, similarly to step 210...

...The procedure repeats a loop through the steps from 218 to 236 inclusive, which either continues mirroring or else, ends mirroring if so commanded.

Claim 23

Provisional Application, written description support for Claim 23:

See page 1, lines 15 - 18:

The description relates to the capability to replicate the contents of a storage device, pertaining to a local SAN (Storage Area Network), in part or in whole, to a remote-SAN coupled to the local SAN replication or mirroring understood as copying, are applicable to any object, volume, file, data structure or the like.

Provisional Application, enablement support for Claim 23:

See page 1, lines 15 - 18:

The description relates to the capability to replicate the contents of a storage device, pertaining to a local SAN (Storage Area Network), in part or in whole, to a remote-SAN coupled to the local SAN replication or mirroring understood as copying, are applicable to any object, volume, file, data structure or the like.

International Application, written description for Claim 23:

See page 1, lines 11 - 12:

A data object is a volume, a logical or virtual volume, a data file, or any data structure. The terms data object and volume are used interchangeably below.

International Application, enablement support for Claim 23:

See page 1, lines 11 - 12:

A data object is a volume, a logical or virtual volume, a data file, or any data structure. The terms data object and volume are used interchangeably below.

Claim 24

Provisional Application, written description support for Claim 24:

See page 2, lines 21 - 24:

The aim of the mirroring functionality is to create the replication of an object. For example, mirroring an object from one storage device to another storage device in the same SAN or remote mirroring of an object residing in a local site SAN (SANL) to any remote site SAN (SANRi), or vice-versa.

See page 6, lines 16 - 19:

Since the configuration of the source volume 1 is intentionally frozen to prevent any change as from the time t0 of the freeze, all the updates pertaining to that source volume 1, thus certainly all the I/O WRITE instructions directed thereto, are routed to the simultaneously created auxiliary volume 1.

Provisional Application, enablement support for Claim 24:

See page 10, line 24 to page 12 line 23:

When mirroring is ordered, as in step 202, for example for a certain volume FL residing at the local site SANL, to be mirrored to a given remote site SRx, then the next step, here 204, is to set a cascade counters to one (s = 1)...

...Since the control loop repeats the run from step 218 to step 236 in Fig. 7, the repetitive stepwise procedure is called cascading.

International Application, written description for Claim 24:

See page 5, lines 14 - 18:

It is yet a further object of the present invention to provide a method and a system for:

storing in the at least one remote storage device of a complete mirrored copy of the selected data object comprising updates entered thereto at the time when copy of the before to penultimate local auxiliary volume was completed.

See page 6, lines 10 - 14:

Mirroring consists of a succession of freeze and copy procedures repeated sequentially in successive mirroring cycles. Only the last local updated mirrored version is saved in the remote storage device. Each new updated version overwrites the previous version. An updated version existing when mirroring starts with a first mirroring cycle $s = 1$, is safely stored after two more mirroring cycles, when $s = 3$.

See page 6, line 23 - 25:

Since the source volume is and must remain frozen, it may not incur changes, but it may be copied by the copy procedure to the remote storage device.

See page 6, lines 32 - 33:

Once the frozen source volume is mirrored to the remote storage device, the freeze and copy procedures are repeated.

See page 6, line 40 to page 7 line 2:

In principle, the mirroring functionality performs successive freeze and copy procedures to replicate one, or a group of data object(s), from one or more local storage device(s), to one or more other, or remote, storage device(s).

See page 8, lines 11-13:

As stated above, the mirroring of a data object from one storage device to another storage device requires the application of successive freeze and copy procedures.

See page 14, lines 11 - 15:

Referring to the description related to the freezing of a source volume SV, at stage 2a in Fig. 2, it was stated that

the source volume was copied after the freeze took place. The mirroring functionality may thus command to copy the frozen source volume SV, from the storage device of origin wherein it resides, defined as a local storage device, to any other storage device, which is referred to as a remote storage device.

International Application, enablement support for Claim 24:

See page 14, line 28 to page 18, line 33:

In Fig. 7 the left column relates to the local storage device SDL wherein a data object resides in the source volume SV, and the abscise displays a time axis t...

...The procedure repeats a loop through the steps from 218 to 236 inclusive, which either continues mirroring or else, ends mirroring if so commanded.

Claim 25

Same as for claim 1.

Claim 26

Same as for claim 2.

Claim 27

Same as for claim 3.

Claim 28

Same as for claim 4.

Claim 29

Same as for claim 5.

Claim 30

Same as for claim 6.

Claim 31

Same as for claim 7.

Claim 32

Same as for claim 8.

Claim 33

Same as for claim 9.

Claim 34

Same as for claim 10.

Claim 35

Same as for claim 11.

Claim 36

Same as for claim 12.

Claim 37

Same as for claim 13.

Claim 38

Same as for claim 14.

Claim 39

Same as for claim 15.

Claim 40

Same as for claim 16.

Claim 41

Same as for claim 17.

Claim 42

Same as for claim 18.

Claim 43

Same as for claim 19.

Claim 44

Same as for claim 20.

Claim 45

Same as for claim 21.

Claim 46

Same as for claim 22.

Claim 47

Same as for claim 23.

Claim 48

Same as for claim 24.

In view of the foregoing, Applicant respectfully states that each claim is supported by the provisional application and the International Application, and that Applicant has been fully responsive to the interrogatories set forth on page 2 of the August 11, 2006 Office Action.

PRIOR ART REJECTIONS

In the Office Action claims 1-3, 5-27 and 29-48 are rejected under 35 USC 102(b) as being anticipated by USP 5,835,953 (Ohran). Claims 4 and 28 are rejected under 35 USC 103(a) as being unpatentable over Ohran in view of U.S. Patent Publication No. 2002/0156971 (Jones). In response, claims are amended and remarks are set forth below in a sincere effort to point out patentable features of the present invention.

The amendments to claims 1 and 25 are supported in the title of the present application, namely "Asynchronous Mirroring in a Storage Area Network," in the Technical Field, on page 1, lines 5-7 and on page 6 lines 5-9 of the application as originally filed, inter alia.

Regarding the Examiner's statement in Item #6 of the Office Action concerning "a mirroring functionality running in the first and in the second processing facility," it is respectfully maintained that the asynchronous mirroring functionality recited in the present claimed invention is not to be confused with the traditional mirroring operation known in the art and referred to in Ohran at column 9, lines 14 - 40 wherein it is stated:

Backup system 14 of FIG. 1 comprises backup storage means for storing data blocks received from primary system 12. Backup storage means can comprise any type of storage device capable of storing blocks of data received from a primary system. For example, backup storage means may comprise a storage device identical to the mass storage device of a primary system. If the primary system has a large magnetic disk, for example, the backup storage means may also

comprise a large magnetic disk. If the backup storage means is the same as the mass storage means of the primary system, the backup storage means can closely mirror the mass storage means of the primary system. As another example, backup storage means may comprise archival storage devices such as a magnetic tape drive or an optical or magneto-optical drive. The type of storage devices that may be used for backup storage means is limited only by the particular application where they are utilized. In some situations it may be more desirable to have a backup storage means that more closely resembles the mass storage means of the primary system. In other situations it may be perfectly acceptable to have archival type storage means that are optimized to store large amounts of data at the expense of rapid access, storage means be able to store the backup storage means be able to store data blocks transferred to the backup system from the mass storage means of the primary system. In FIG. 1 the backup storage means is illustrated by backup storage device 24.

In contrast with the teachings of Ohran, the mirroring functionality recited in the present claimed invention is different at least by being an asynchronous mirroring procedure. The mirroring functionality first freezes a selected data object, then copies the frozen selected data object into at least one remote storage device, and allows parallel continued use and updating of the selected data even with the frozen object, and commands unconditional repeated run by default as long as desired, as recited in Claim 1.

In contrast with the mirroring functionality of the present claimed invention, Ohran teaches a version of a backup system taking snapshots, as recited in the title of the patent.

Ohran teaches "a backup system that takes a snapshot" and not "a mirroring system.". Ohran was well aware of the

limitations of traditional disk mirroring techniques, and recognized them as enabling to make only complete storage copies, as recited in column 5, lines 39 - 57:

The present invention begins with the assumption that a primary mass storage device connected to a primary computer and a backup storage system connected to a backup storage device contain identical data. This may be accomplished, for example, by making a complete copy of the primary mass storage device to the backup device using either traditional backup techniques or traditional disk mirroring techniques. Once the primary mass storage device and the backup storage device contain the same data, the present invention tracks the changes made to the primary mass storage device. This tracking is done by identifying those storage locations on the primary mass storage device that have new data written in them from the time that the backup storage device was in sync with the primary mass storage device. By identifying those changes that have been made to the primary mass storage device, the invention identifies those changes that need to be made to the backup storage device in order to bring the backup storage device current with the primary mass storage device.

Regarding "default backup operations," Ohran recites at column 6, lines 49 - 52, that his backup system cannot run by unconditioned default, but is limited to capture only logically consistent and specific states:

Yet another object of the present invention is to allow the backup system to capture successive logically consistent backup states in order to provide a series of logically consistent backup states.

Ohran teaches a backup system and method that takes snapshots of data only on predetermined conditions, such as the data residing in a logically consistent state, as recited at column 5, lines 20-38:

The foregoing problems in the prior state of the art have

been successfully overcome by the present invention, which is directed to a system and method for backing up a primary mass storage device to a backup storage device. The current system and method provides three significant advantages over the prior art. First, the backup system and method of the present invention reduces the amount of data needed to make a backup by backing up only those data blocks of the primary mass storage device that change. Second, the system and method of the present invention emphasize security of the backup by ensuring that the primary storage device is in a logically consistent state when a backup is made. Third, because the data needed to make a backup is reduced to the absolute minimum, and because backups are only made of logically consistent states, backup frequency can be increased in order to capture many more logically consistent states. The backup system and method of the present invention accomplishes this without terminating user access to the mass storage system.

The limitations indicated by Ohran are necessary for backup systems, for example when there is a need to restore data. Should inconsistent files reside in the backup, then, at the time a restoration is requested, those inconsistent files will corrupt the current files residing in the primary system. This may happen even if the files in the primary system reside in a logical consistent version.

In contrast to Ohran, with the asynchronous mirroring functionality of the present claimed invention, the mirrored data at the remote site is permitted to contain logical inconsistencies, which is different from Ohran. With the present claimed invention there is no need for logical consistency as a prerequisite for mirroring or for restoration since the recovery is achieved by simply mounting the file system of the remote mirroring functionality. Should logical inconsistencies be

present at recovery time, these inconsistencies will be treated at the time when the files system is mounted in the same manner as a file system recovery process would treat a sudden lack of power when in the midst of operations.

Instead of "commanding, by default, repeated run of the mirroring functionality for copying updates to the selected data object, unless receiving command for mirroring break" as recited in claim 1 of the present claimed invention, Ohran depends on logical state conditions being achieved. In other words, Ohran may take a snapshot and backup data but only on provision that logical state conditions are met, in contrast with the present claimed invention. With Ohran such a backup command is not applicable when desired, but only when logical consistency conditions are met.

It should be appreciated that because no logical state conditions are required with the present claimed invention, the frequency of mirroring replication may be increased in comparison to Ohran. A logical consistent condition is achieved by requiring, or by waiting until the file system stops all of the I/O and write operations of its cached data into storage, which may take a few seconds. While it is reasonable to wait for a logical consistent condition and to mirror every few hours, it is obvious that stopping for a few seconds cannot be permitted every few minutes. Due to the fact that the present claimed invention

does not require stopping or waiting for logical consistent conditions, asynchronous functionality mirroring operations may be more frequent, such as every few minutes. Thereby a much more granular recovery may be achieved. Even if taken alone, this feature presents an enhancement of orders of magnitude when compared to Ohran.

As stated above, and in contrast with the present claimed invention, Ohran takes snapshots. The static snapshots taken by Ohran are independent entities that are first stored in a snapshot storage device, as recited at column 8, lines 51-65 of Ohran:

As described in greater detail below, embodiments within the scope of this invention use a static snapshot of all or part of the mass storage device during the backup process. Embodiments within the scope of this invention therefore comprise preservation memory means for storing data blocks of said mass storage means so as to create a static snapshot of the mass storage means at a particular point in time. As described in greater detail below, such preservation memory means may comprise any type of writable storage device such as RAM, EEPROM, magnetic disk storage, and the like. Such preservation memory means may also comprise a portion of mass storage device 20. In FIG. 1, such preservation memory means is illustrated, for example, by snapshot storage device 22. Preservation memory means is discussed in greater detail below.

Hence, Ohran teaches a backup system and method that takes a snapshot only when the data resides in a logically consistent state, which data is then stored in a snapshot storage device 22. The snapshot storage device 22 is an intermediary station wherein a copy of the data to be backed up is saved prior to transfer

into backup storage, as recited by Ohran at column 8, lines 3-25:

Referring now to FIG. 1, a system level block diagram of one embodiment of the present invention is presented. The system, shown generally as 10, comprises one or more primary systems 12, a backup system 14, and backup transport means for transporting data between primary system 12 and backup system 14. In FIG. 1, the backup transport means is illustrated as backup transport link 16. In FIG. 1, primary system 12 may be any type of networked or stand alone computer system. For example, primary system 12 may be a network server computer connected to a computer network such as computer network 18. Primary system 12 may also be a stand alone system. Primary system 12 may also be a backup or standby server of a computer network connected to a primary server. The present invention can be used with any type of computer system. In this sense, the term "primary" is not meant to define or describe a computer system as a primary network server (as opposed to a backup or standby network server). In this description, the term "primary" is used to refer to the fact that the system has attached mass storage means for storing a copy of the data that is to be backed up. In other words, the term "primary" is used to differentiate the system from backup system 14.

For backing up snapshot data saved in a snapshot storage device 22 residing in the primary system 12, Ohran recites a backup storage device 24 at column 9, lines 14-40:

Backup system 14 of FIG. 1 comprises backup storage means for storing data blocks received from primary system 12. Backup storage means can comprise any type of storage device capable of storing blocks of data received from a primary system. For example, backup storage means may comprise a storage device identical to the mass storage device of a primary system. If the primary system has a large magnetic disk, for example, the backup storage means may also comprise a large magnetic disk. If the backup storage means is the same as the mass storage means of the primary system, the backup storage means can closely mirror the mass storage means of the primary system. As another example, backup storage means may comprise archival storage devices such as a magnetic tape drive or an optical or magneto-optical drive. The type of storage devices that may be used for backup storage means is limited only by the particular application where they are utilized. In some situations it

may be more desirable to have a backup storage means that more closely resembles the mass storage means of the primary system. In other situations it may be perfectly acceptable to have archival type storage means that are optimized to store large amounts of data at the expense of rapid access, storage means be able to store the backup storage means be able to store data blocks transferred to the backup system from the mass storage means of the primary system. In FIG. 1 the backup storage means is illustrated by backup storage device 24.

It is noted that among other options, Ohran also recites the use of traditional mirroring for copying the contents of the snapshot storage device 22 to the backup storage device 24.

In conclusion, Ohran takes a snapshot of data, makes a copy of the snapshot for storage into a snapshot storage device, and finally transfers copies of the snapshot from the snapshot storage device into a backup storage device.

Ohran does not disclose, teach or suggest the use of "a mirroring functionality" having "a freeze procedure for freezing the selected data object," but rather takes a snapshot. Furthermore, Ohran does not disclose, teach or suggest implementing "a copy procedure for copying the frozen selected data object into the at least one remote storage device" but instead stores a snapshot, taken only under predetermined conditions, into a snapshot storage and copies the snapshot out of snapshot storage into a backup storage. The required preconditions, i.e., logical states, prevent Ohran from making backups by default whenever desired.

Ohran and the present claimed invention are different in

concept and in embodiment. As compared to Ohran, the present claimed invention is a "forward snapshot," while Ohran uses a "backward snapshot." This means that the present claimed invention freezes the current data set and writes "new data," or updates received data after a freeze operation, into an auxiliary volume while with Ohran the "new data" is written into the primary system, and the "old data" is saved as a snapshot in the mass storage as a copy. This difference permits the following abilities, not recited by, and not available with Ohran:

- * to keep multiple point-in-time copies in the source volume, thereby allowing the recovery of data without the need to fetch that data from a secondary location;

- * to keep multiple point-in-time copies in the destination, or remote location, thereby saving previous versions and not just the last version;

- * to allow multiple independent hosts, or machines, or users, to write to copies of the frozen data, since each host writes data into its own auxiliary volume without interfering with other hosts. Since with Ohran updates are directed to the the primary system volume, then if two hosts want to write "new data" after a snapshot, the data of each one will corrupt the data of the other hosts;

- * to retrieve one of the point-in-time versions by help of the present claimed invention, it suffices to point to the source

and add the additional increments, which is an instant, real time operation. With Ohran, incremental data has to be applied back to the primary source volume, but in the mean time the source has already collected further updates, and therefore needs to be rolled back in time to apply the increments to the old data. Evidently, the need to rollback has a negative time duration influence on the recovery time. While the present claimed invention permits recovery in seconds, Ohran will need to wait until all of the increments are applied back. Furthermore, with the present claimed invention, if the roll back operations eventually fails, nothing is lost, as it is still possible to refer to a different point in time and move the time forward, because data does not need to be transferred. In Ohran, due to the fact that the older data is applied on top of the newer data, should a recovery operation fail, it is not possible to refer to another point in time since the older data erased the new data during the first roll back operation.

It is respectfully submitted that the present claimed invention is for asynchronous mirroring to at least one remote storage device and is different from Ohran, since Ohran does not disclose, teach or suggest freezing of a selected data object but the taking of a snapshot of the selected data object and thereafter storage of the snapshot into a backup storage. In contrast to Ohran, the present claimed invention freezes the

selected data object for mirroring to at least one remote storage device, does not take snapshots that are therefore not stored in memory, and is not limited to mirroring only in logically consistent states. It is therefore submitted that the incremental "backward snapshot" technology used by Ohran is different from the present claimed invention that recites asynchronous mirroring, "forward operative" freeze of data, and making copies of the frozen data to remote backup storage, without requiring any snapshot and any snapshot storage device.

In view of the foregoing, claim 1 is patentable over Ohran under 35 USC 102 as well as 35 USC 103.

None of the other references of record close the gap between the present claimed invention as defined by claim 1 and Ohran. Therefore, claim 1 is patentable over all of the references of record when taken either alone under 35 USC 102 or in combination under 35 USC 103.

Claim 25 is patentable over Ohran and all of the references of record for reasons, inter alia, set forth above in connection with claim 1.

Claims 2, 3, 5-24, 26, 27 and 29-48 are either directly or indirectly dependent on claims 1 or 25, and are patentable over the cited references in view of their dependence on claims 1 or 25, and because the references do not disclose, teach or suggest each of the limitations set forth in the dependent claims.

With respect to the dependent claims, Applicant has the following additional comments.

In item #7 of the Office Action, the Examiner states:

"With respect to claims 2 and 26, the mirroring functionality further comprises:

the freeze procedure being applied for freezing the selected data objects as a source volume (SV) (figure 3, 22, 48), at least one local auxiliary volume (AVL) to which updates addressed to the selected data are redirected (20 + 52), each single data object out of the selected data object corresponding to one local auxiliary volume out of the at least one local auxiliary volume, at least one remote volume being created in each remote storage device out of the at least one remote storage device, to correspond to each one local auxiliary volume created, a resulting source volume being formed in the at least one local storage device to comprise the frozen selected data object and the at least one auxiliary volume, and the copy procedure being applied for copying the frozen selected data object from the at least one resulting volume into the at least one remote storage device (figure 7B, see also col. 5, line 39 - col. 6, line 18)."

Regarding "the freeze procedure being applied for freezing the selected data objects as a source volume (SV) (figure 3, 22, 48)," it is remarked that Ohran does not teach the use of freezing or a freeze procedure and that the words freeze or freezing are not recited in Ohran. Furthermore, Ohran does not disclose the creation of a source volume SV. Even though Ohran recites the taking of a snapshot of selected data, that snapshot is not frozen as a source volume (SV) but rather stored in a snapshot storage device, as recited by Ohran at column 8, lines 51 - 65:

As described in greater detail below, embodiments within the scope of this invention use a static snapshot of all or part of the mass storage device during the backup process. Embodiments within the scope of this invention therefore comprise preservation memory means for storing data blocks of said mass storage means so as to create a static snapshot of the mass storage means at a particular point in time. As described in greater detail below, such preservation memory means may comprise any type of writable storage device such as RAM, EEPROM, magnetic disk storage, and the like. Such preservation memory means may also comprise a portion of mass storage device 20. In FIG. 1, such preservation memory means is illustrated, for example, by snapshot storage device 22. Preservation memory means is discussed in greater detail below.

Ohran also recites the snapshot storage device in relation to Fig. 3 from column 13, line 63 to column 14, line 7:

In FIG. 3, the means for preserving a static snapshot is illustrated by snapshot processing block 50. As illustrated in FIG. 3, it may make sense to incorporate the snapshot processing mechanism into the mass storage read/write processing block. Although the details of snapshot processing block 50 are presented below, one preferred embodiment preserves a static snapshot by copying a data block that is to be overwritten from mass storage device 20 into snapshot storage 22 and then indicating that the block has been preserved in snapshot map 52. Once a copy has been placed into snapshot storage 22, then the copy of the data block on mass storage device 20 can be overwritten.

In figure 3, item 22 is the snapshot storage device and item 48 is a backup map of the snapshot storage device, both items 22 and 48 being related to snapshot storage but without relation to freezing of the selected data objects as a source volume (SV) that is recited in claim 2, but is not taught by Ohran.

Reference is now made to "at least one local auxiliary volume (AVL) to which updates addressed to the selected data

object are redirected (20 + 52), each single data object out of the selected data object corresponding to one local auxiliary volume out of the at least one local auxiliary volume." Ohran does not create an auxiliary volume to which updates are directed but continues to send updates to the selected data after having taken a snapshot. For data blocks that are to be preserved, first a snapshot is taken, which is copied into the snapshot storage and new data may then be written to the data blocks residing in the mass storage device 20 without losing the preserved data blocks. Thereby, selected data blocks may be updated or written to after taking the snapshot, as recited by Ohran at column 17, lines 33-55:

Returning now to FIG. 4, if a snapshot request has been received, then the next step is to preserve a static snapshot of at least a portion of mass storage device 20. Although any means to preserve a static snapshot can be used with the present invention, it is preferred that a particular process be used to preserve a static snapshot. The preferred method is summarized in the description of steps 70, 72, 74, decision block 84, and step 86 below. The method is more particularly described in U.S. Pat. No. 5,649,152 entitled METHOD AND SYSTEM FOR PROVIDING A STATIC SNAPSHOT OF DATA STORED ON A MASS STORAGE SYSTEM, previously incorporated by reference. In essence, a preferred method of preserving a static snapshot utilizes a snapshot storage, such as snapshot storage 22 of FIG. 3, to preserve data blocks of a mass storage device, such as mass storage device 20 of FIG. 3, that are to be overwritten with new data. As explained in greater detail below, the data blocks that are to be preserved are first copied into the snapshot storage and a record indicating that the data block has been preserved is updated. Such a record can be stored, for example, in snapshot map 52 of FIG. 3. New data may then be written to mass storage device 20 without losing the preserved data blocks.

It is assumed that "(20 + 52)" are items related to Fig. 3, where in the primary system 12, the index 20 indicates the mass storage and 52 marks the snapshot map. The statement that "updates addressed to the selected data object are redirected (20 + 52)" is not applicable with Ohran. In claim 2 of the present claimed invention, a local source volume SV, containing an auxiliary volume AVL is created, and it is to the auxiliary volume AVL that updates addresses to the selected data object are redirected.

With Ohran, updates addressed to the selected data blocks reach their address, thus the selected data blocks, and are therefore not redirected. It is noted that Ohran does not create auxiliary volumes.

Claim 2 of the present claimed invention also recites that a local auxiliary volume becomes one portion out of a resulting source volume that includes both the frozen selected data object and the at least one local auxiliary volume. Claim 12 of the present claimed invention recites in more detail that updates addressed to the selected data object are redirected to the local auxiliary volume.

In contrast to the present claimed invention, Ohran takes a snapshot but directs updates to the selected block of data, but not to an auxiliary volume that is absent from his disclosure. Ohran teaches in column 17, lines 33-55 that selected data blocks

that are at least a portion of mass storage device 20 are preserved in a snapshot so that updates with new data may be written to mass storage device 20 without losing the preserved data blocks that have previously been saved by taking a snapshot (see Ohran at col. 17, lines 33-55).

With Ohran "updates addressed to the selected data object" are not redirected to anywhere but to the selected data since no auxiliary volume is created for that purpose as is the case with the present claimed invention. Therefore Ohran's recitation does not disclose, teach or suggest "each single data object out of the selected data object corresponding to one local auxiliary volume out of the at least one local auxiliary volume" because there are no auxiliary volumes.

In response to the statement "at least one remote volume being created in each remote storage device out of the at least one remote storage device, to correspond to each one local auxiliary volume created," it was indicated hereinabove that Ohran does not implement local auxiliary volumes in primary storage. Furthermore, Ohran does not implement auxiliary volumes in backup storage but simply copies data blocks from primary storage to backup storage without any auxiliary volume, as recited at column 9, lines 14-40:

Backup system 14 of FIG. 1 comprises backup storage means for storing data blocks received from primary system 12. Backup storage means can comprise any type of storage device capable of storing blocks of data received from a primary

system. For example, backup storage means may comprise a storage device identical to the mass storage device of a primary system. If the primary system has a large magnetic disk, for example, the backup storage means may also comprise a large magnetic disk. If the backup storage means is the same as the mass storage means of the primary system, the backup storage means can closely mirror the mass storage means of the primary system. As another example, backup storage means may comprise archival storage devices such as a magnetic tape drive or an optical or magneto-optical drive. The type of storage devices that may be used for backup storage means is limited only by the particular application where they are utilized. In some situations it may be more desirable to have a backup storage means that more closely resembles the mass storage means of the primary system. In other situations it may be perfectly acceptable to have archival type storage means that are optimized to store large amounts of data at the expense of rapid access, storage means be able to store the backup storage means be able to store data blocks transferred to the backup system from the mass storage means of the primary system. In FIG. 1 the backup storage means is illustrated by backup storage device 24.

It is noted that Ohran does not need, and therefore does not mention auxiliary volumes, neither in the primary system 12, nor in the backup system 14. As a result of the fact that Ohran does not teach auxiliary volumes, the statement "a resulting source volume being formed in the at least one local storage device to comprise the frozen selected data object and the at least one auxiliary volume" is not valid with reference to Ohran. Ohran performs backup of a snapshot stored in primary storage into backup storage, which is different from the present claimed invention that mirrors a frozen data block into a remote storage, as recited in claim 12.

Reference is made to the statement "the copy procedure being

applied for copying the frozen selected data object from the at least one resulting volume into the at least one remote storage device (figure 7B, see also col. 5, line 39 - col. 6, line 18)" in Item #7 of the Office Action.

First, it was pointed out above that Ohran does not freeze selected data but takes a snapshot of the selected data that is stored in the primary system and later copied to the backup system. Second, as already mentioned above, Ohran does not recite any resulting volume simply because such a resulting volume includes both frozen selected data and a local auxiliary volume, both of which are not recited and not implemented by Ohran.

Referring to "(figure 7B, see also col. 5, line 39 - col. 6, line 18)," Ohran discloses his backup system that stores a snapshot on condition that a logically consistent state has been reached, as recited from column 5, line 39 to column 6, line 18:

The present invention begins with the assumption that a primary mass storage device connected to a primary computer and a backup storage system connected to a backup storage device contain identical data. This may be accomplished, for example, by making a complete copy of the primary mass storage device to the backup device using either traditional backup techniques or traditional disk mirroring techniques. Once the primary mass storage device and the backup storage device contain the same data, the present invention tracks the changes made to the primary mass storage device. This tracking is done by identifying those storage locations on the primary mass storage device that have new data written in them from the time that the backup storage device was in sync with the primary mass storage device. By identifying those changes that have been made to the primary mass storage device, the invention identifies those changes that

need to be made to the backup storage device in order to bring the backup storage device current with the primary mass storage device.

Once the changes that need to be made to the backup storage device have been identified, the changes are sent to the backup system. The backup system then has available all data to bring the backup storage device current with the primary mass storage device. In order to preserve the original data of the primary mass storage device during the backup process, a static snapshot of the primary mass storage device is taken. This static snapshot captures the changes that have been made to the primary mass storage device and that need to be transferred to the backup system. In order to make the backup transparent to users, it is preferred that the static snapshot be taken in such a way that user access to the primary mass storage device is not interrupted.

The present invention includes a mechanism to identify when the primary mass storage device is in a logically consistent state in order to determine when a static snapshot should be made. By identifying a logically consistent state and the taking a static snapshot of the changes made up to that point in time, when the changes are transferred to the backup system, the backup system is guaranteed to capture a logically consistent state. By capturing snapshots of succeeding logically consistent states, the backup can capture one logically consistent state after another. In this way, if the backup data should ever be needed, the backup data will be in a logically consistent state. The backup system moves from one logically consistent state to another logically consistent state thus eliminating one of the problems of the prior art.

As indicated above, there is no "copy procedure being applied for copying the frozen selected data object from the at least one resulting volume into the at least one remote storage device" as Ohran recites neither a frozen selected data object nor an at least one resulting volume.

With respect to "figure 7B," Ohran provides details about the creation of a backup from mass storage residing in the

primary system to the backup storage located in the backup system, as already referred to hereinabove, as recited in column 22, from line 39 to line 52:

In order to illustrate in greater detail the operation of FIGS. 3-6 in creating a backup, a detailed example is presented in FIGS. 7A and 7B. Referring first to FIG. 7A, consider a group of data blocks 122, stored in storage locations numbered 1-6, of mass storage device 20. FIG. 7B shows that backup storage device 24 also has a similar group of data blocks 124, also stored in storage locations numbered 1-6. At time T.sub.0, the data blocks stored in 122 are identical to the data blocks stored in 124. Referring again to FIG. 7A, backup map 48 has six map locations 126 that correspond to storage locations 122. Snapshot map 52 also has six map locations 128 that correspond to storage locations 122. As illustrated in FIG. 7A, at time T₀ map location 126 and 128 are cleared.

Further reference to Fig. 7B is also made by Ohran in column 25, lines 3 to 19, where mention is made of a backup capture buffer 26, inexistent and unnecessary with the present claimed invention, which implements mirroring in contrast to transfer from primary storage to backup storage as does Ohran:

As described in greater detail below, when data blocks are received by a backup system it may be desirable to store the data blocks as they are received in a backup capture buffer, such as backup capture buffer 26 of FIG. 3. This allows all data blocks to be received before they are applied to backup storage device 24 or before they are saved as an incremental backup. In FIG. 7B, data blocks 152 are received by the backup system and applied to storage locations 124 to achieve storage locations 154. Storage locations 154 are identical to storage locations 138 of the primary system (FIG. 7A). Recall that storage locations 138 represented the state of mass storage device 20 at time T.sub.1 when the snapshot was taken. Thus, the changes that have occurred between time T.sub.0 and time T.sub.1 have now been backed up to the backup system and applied to backup storage device 24 in order to bring backup storage device 24 current with mass storage device 20 at time T.sub.1.

Following the aforementioned arguments, it is respectfully maintained that the present claimed invention is patentable over Ohran. The reasons provided hereinabove with respect to claim 2 are also valid for claim 26, which is a corresponding system claim.

According to the arguments provided above, Ohran operates a backup system that is different from the present claimed invention, which is directed to a mirroring functionality. It is therefore respectfully requested that the rejection of claims 2 and 26 be withdrawn.

In item #11 of the Office Action the Examiner rejects claims 4 and 28 under 35 U.S.C. 103 as being unpatentable over Ohran in view of Jones (U.S. Patent Application Publication No. 2002/0156971, hereinafter "Jones").

Claims 4 and 28 are dependent on claims 1 and 25 respectively, and are patentable over the cited references in view of their dependence on claims 1 or 25 and because the references do not disclose, teach or suggest each of the limitations set forth in claims 4 and 28.

In view of all of the foregoing, claims 1-48 are in form for immediate allowance, which action is earnestly solicited.

FORM PTO-892

The Examiner is respectfully advised that he has applied

U.S. Patent Application Publication No. 2002/0156971 (Jones)
against claims 4 and 28. The Examiner is respectfully advised
that the Jones reference is not identified on the Form PTO-892
attached to the August 11, 2006 Office Action, and the reference
has not been previously cited by the Applicant in Applicant's
February 10, 2004 Information Disclosure Statement. In order to
ensure that the Jones reference is made of record in connection
with the application, the Examiner is respectfully requested to
issue a Form PTO-892 listing the Jones reference thereon.


* * * * *

Entry of this Amendment, allowance of the claims and the
passing of this application to issue are respectfully solicited.

If the Examiner disagrees with any of the foregoing, the
Examiner is respectfully requested to point out where there is
support for a contrary view.

If the Examiner has any comments, questions, objections or recommendations, the Examiner is invited to telephone the undersigned at the telephone number given below for prompt action.

Respectfully submitted,



Robert P. Michal
Reg. No. 35,614

Frishauf, Holtz, Goodman & Chick, P.C.
220 Fifth Avenue
New York, New York 10001-7708
Tel. (212) 319-4900
Fax (212) 319-5101
RPM/ms

Encl: Petition For Extension of Time